NTPv4 Protocol Specification

Jack L. Burbank (Editor), Jim Martin (Editor), David Mills

jack.burbank@jhuapl.edu

Presentation to NTP WG
NTP 63 – Paris, France

Document Overview

- <draft-ietf-ntp-ntpv4-proto-00.txt> posted on 11 July 2005
- Document to be published as Proposed Standard
- Goal is to have document in WG last call by IETF 64 (November 05)

Document Overview (continued)

- Approach to writing protocol specification:
  1. Compile material ‘as-is’ from existing material, putting it into proper format
     - RFC 2030 (SNTPv4 specification)
     - draft-mills-sntp-v4-00.txt (update to RFC 2030 currently in RFC Editor Queue)
     - RFC 1305 (NTPv3 specification)
     - Additional material from NTP.org (e.g. flow.ppt, secproto.ppt)
  2. Add any additional material that is required
  3. Clean up issues as necessary
     - Outstanding IESG comments
       - https://datatracker.ietf.org/public/pidtracker.cgi?command=view_id&dTag=10845&rfc_flag=0

IESG Comments (General)

- “This document has some significant issues/errors regarding IPv6 addressing and anycast…”
- “…document is confused about the use of anycast, multicast and unicast addresses…”
- “…I believe the anycast language needs serious work, as does the kiss of death. I fervently hope that the entire best practices section will be re-written from the ground up, as it could now be read to imply that backoff should stop at 1 minute poll intervals.”
- “The Security Considerations section is unacceptable. Given the fact that an authentication mechanism is defined, I think it is straightforward to add a paragraph about the consequences of not using the authentication mechanism…”

Summary of Detailed IESG Comments

- Several specific comments regarding usage of multicast/anycast/broadcast terminology
- Need to add IPv6 reference
- Need to add MLD reference
- How would cryptographic authentication apply to broadcast/multicast/anycast services?
- How does client/server behave if messages with invalid field values are received?
- Removal of implementation details
- Clarify text regarding polling interval

Other Comments Received to Date…

- Need to resolve reference to IPv6 site-locals
- Need to add reference to MLDv2

Additional Changes

- Description of each mode of operation in “Protocol Operation” section?
  - Or reference architecture draft?
- Addition of terminology section?
  - Or reference architecture draft?
- Explicit breakout of SNTPv4 and NTPv4 in document?
- Detailed description of optional extension fields
- Expansion of tables in Sections 5 and 6 (client operations and server operations, respectively) to include all modes of operation
- Additional detail on autokey protocol
  - Draw material from “The Autokey Security Architecture, Protocol and Algorithms,” by D. Mills
- IANA considerations?
- Security considerations?

Way Forward

- New draft to be distributed within 3-4 weeks of IETF-63
  - Goal to address IESG comments and any additional comments received from working group
- Strongly encourage everyone to review and provide comments on the mailing list
- One last thing to think about…
  - Advancement beyond proposed standard requires multiple ***independent*** implementations

Backup Slides

IESG Comments (Detailed – 1/11)

- 2. Operating Modes and Addressing

  Unless excepted in context, reference to broadcast address means IPv4 broadcast address, IPv4 multicast group address or IPv6 site-local scope address.

  >> So, a "broadcast address" may be an IPv6 site-local (unicast)
  >> address? This seems particularly obscure for two reasons: (1)
  ...snip...
  >>... broadcast addresses. Is
  >> there any reason for this, other than a lack of desire to change
  >> the text of the original document?

- Further information on the broadcast/multicast model is in RFC-1112 [DEE89]. Details of address format, scoping rules, etc., are beyond the scope of this memo. SNTPv4 can operate with either unicast (point to point), broadcast (point to multipoint) or anycast (multipoint to point) addressing modes.

  >> Anycast addresses are not "multipoint to point". They are not
  >> easily mapped onto this terminology, since they are a special type
  >> of "point to multipoint" address that can be used to establish
  >> "point to nearest point" communication.

IESG Comments (Detailed – 2/11)

- Anycast is designed for use with a set of cooperating servers whose addresses are not known beforehand. The anycast client sends an ordinary NTP client request to a designated broadcast address.

  >> An anycast address shouldn't be a broadcast address, at least not
  >> in IPv6 anycast.

- One or more anycast servers listen on that address. Upon receiving a request, an anycast server sends an ordinary NTP server reply to the client.

  >> This description should indicate what source address the server
  >> uses in this reply. Presumably one of its own normal unicast
  >> addresses? Does it need to be the same IP version (v4 or v6) as
  >> the anycast address on which the request was received?

IESG Comments (Detailed – 3/11)

- The client and server addresses are assigned following the usual IPv4, IPv6 or OSI conventions. For NTP multicast, the IANA has reserved the IPv4 group address 224.0.1.1 and the IPv6 group address ending :101, with prefix determined by scoping rules. The NTP broadcast address for OSI has yet to be determined. Notwithstanding the IANA reserved addresses, other multicast addresses can be used which do not conflict with others assigned in scope. In the case of IPv4 multicast or IPv6 broadcast addresses, the client must implement

  >> There is no such thing as an IPv6 broadcast address.

- the Internet Group Management Protocol (IGMP) as described in RFC-3376 [CAIN02], in order that the local router joins the multicast

  >> A reference to MLD (RFC 3810) should be included here.

IESG Comments (Detailed – 4/11)

- In the case of SNTP as specified herein, there is a very real vulnerability that SNTP broadcast clients can be disrupted by misbehaving or hostile SNTP or NTP broadcast servers elsewhere in the Internet.

  >> It is not my understanding that broadcasts (of any type) are
  >> typically forwarded across the Internet, or even between local
  >> subnets. Perhaps this is intended to discuss multicast servers?

- It is strongly recommended that access controls and/or cryptographic authentication means be provided for additional security in such cases.

  >> How would cryptographic authentication apply to a broadcast,
  >> mutlicast or anycast service?

IESG Comments (Detailed – 5/11)

- While not integral to the SNTP specification, it is intended that IP broadcast addresses will be used primarily in IP subnets and LAN segments including a fully functional NTP server with a number of dependent SNTP broadcast clients on the same subnet, while IP multicast group addresses will be used only in cases where the TTL is engineered specifically for each service domain.

  >> The discussion of TTL use for this purpose should either include a
  >> reference or be described in more detail in the eventual
  >> specification.

- 4. Message Format

  Both NTP and SNTP are clients of the User Datagram Protocol (UDP) specified in RFC-768 [POS80], which itself is a client of the Internet Protocol (IP) specified in RFC-791) [DAR81].

  >> s/client of/application of/ ??
  >> Also, a reference to IPv6 seems to be missing.

IESG Comments (Detailed – 6/11)

- The structure of the IP and UDP headers is described in the cited specification documents and will not be detailed further here. The UDP port number assigned by the IANA to NTP is 123. The SNTP client should use this value in the UDP Destination Port field for client request messages. The Source Port field of these messages can be any nonzero value chosen for identification or multiplexing purposes. The server interchanges these fields for the corresponding reply messages.

  This differs from the RFC-2030 specifications which required both the source and destination ports to be 123. The intent of this change is to allow the identification of particular client implementations (which are now allowed to use unreserved port numbers, including ones of their choosing)

  >> In what way does allowing the client port to be unspecified allow
  >> the "identification of particular client implementations"? The
  ...snip...
  >>... single system (sharing a single IP
  >> address), which seems to be missed altogether here. Would that
  >> even apply to SNTP?

- Leap Indicator (LI): This is a two-bit code warning of an impending leap second to be inserted/deleted in the last minute of the current day. This field is significant only in server messages, where the values are defined as follows:

  >> The values aren't defined until after the packet format, which is
  >> confusing. This is probably just a formatting error?

IESG Comments (Detailed – 7/11)

-

LI	Meaning
0	no warning
1	last minute has 61 seconds
2	last minute has 59 seconds)
3	alarm condition (clock not synchronized)

>> The alarm condition of the Leap Indicator seems important. What
>> are clients supposed to do if they receive this flag in a response?
>> The later text describes when servers should set it (at start-up,
>> before finding a time source), but doesn't say what a client should
>> do if this flag is received.

- Mode: This is a three-bit number indicating the protocol mode. The values are defined as follows:

Mode	Meaning
0	reserved
1	symmetric active
2	symmetric passive
3	client
4	server
5	broadcast
6	reserved for NTP control message
7	reserved for private use

In unicast and anycast modes, the client sets this field to 3 (client) in the request and the server sets it to 4 (server) in the reply. In broadcast mode, the server sets this field to 5 (broadcast). The other modes are not used by SNTP servers and clients.

>> What should SNTP servers and clients do if they receive messages
>> with other modes set? There is ambiguous/conflicting text about
>> this later in the document.

IESG Comments (Detailed – 8/11)

- 5. SNTP Client Operations

  A SNTP client can operate in unicast, broadcast or anycast modes. In unicast mode the client sends a request (NTP mode 3) to a designated unicast server and expects a reply (NTP mode 4) from that server. In broadcast client mode it sends no request and waits for a broadcast (NTP mode 5) from one or more broadcast servers. In anycast mode, the client sends a request (NTP mode 3) to a designated broadcast

  >> s/broadcast/anycast

- than the selection of address in the request, the operations of anycast and unicast clients are identical.

  >> Will anycast clients retry using the anycast address if their
  >> initial unicast server stops responding?

IESG Comments (Detailed – 9/11)

- 1. When the IP source and destination addresses are available for the client request, they should match the interchanged addresses in the server reply.

  >> This only applies to unicast requests, at least for the client
  >> destination address.

- 2. When the UDP source and destination ports are available for the client request, they should match the interchanged ports in the server reply.

  >> I don't think that these are rational checks at the SNTP level. If
  >> the client address or port doesn't match, the response won't go to
  >> the right host/process. And, checking the server address and port
  >> doesn't seem to have any real benefits.

- 4. The server reply should be discarded if any of the LI, Stratum, or Transmit Timestamp fields are 0 or the Mode field is not 4 (unicast) or 5 (broadcast).

  >> This seems to be a section on optional checks, but the value of the
  >> mode field would seem like an important thing to check on any
  >> reply, to me.
  >>
  >> Should the response also be discarded if the LI is set to 3
  >> (unsychronized)?

IESG Comments (Detailed – 10/11)

- the NTP header, and sends a reply (NTP mode 4), possibly using the same message buffer as the request.

  >> The comment about "possibly using the same message buffer" would
  >> seem to be an implementation detail.

- A anycast server listens on the designated broadcast address, but uses its own unicast IP address in the source address field of the reply. Other than the selection of address in the reply, the operations of anycast and unicast servers are identical. Broadcast messages are normally sent at poll intervals from 64 s to 1024 s, depending on the expected frequency tolerance of the client clocks and the required accuracy.

  >> s/poll intervals/intervals

- 7. Configuration and Management

  Initial setup for SNTP servers and clients can be done using a web client, if available, or a serial port if not.

  >> Implementation details, IMO.

IESG Comments (Detailed – 11/11)

- Broadcast servers and anycast clients must be provided with the TTL and local broadcast or multicast group address. Unicast and anycast servers and broadcast clients may be configured with a list of address-mask pairs for access control, so that only those clients or servers known to be trusted will be accepted. Multicast servers and clients must implement the IGMP protocol and be provided with the

  >> s/IGMP/IGMP or MLD

- 1. A client MUST NOT under any conditions use a poll interval less than one minute.

  2. A client SHOULD increase the poll interval using exponential backoff as performance permits and especially if the server does not respond within a reasonable time.

  >> I don't understand what this means... Should the client increase
  >> the poll interval in normal operations? Or only when packets are
  >> lost?